GENERALIZING SELF-NORMALIZED IMPORTANCE SAMPLING WITH COUPLINGS

NICOLA BRANCHINI

An ubiquitous problem in statistics and machine learning is the estimation of expectations involving PDFs with intractable normalising constants. The self-normalised IS (SNIS) estimator, which normalizes the weights, has become the standard approach due to its simplicity. However, most of the state-of-the-art AIS methods adapt the proposal as if the weights had not been normalized. Further, the SNIS has been shown to exhibit high variance in challenging estimation problems, e.g, involving rare events or posterior predictive distributions in Bayesian statistics.

We propose a framework that considers the original task as estimation of a ratio of two integrals. We obtain samples from a joint proposal distribution in an extended space, with two of its marginals playing the role of proposals used to estimate each integral. Importantly, the framework allows us to induce and control a dependency between both estimators. We propose a construction of the joint proposal that decomposes in two (multivariate) marginals and a coupling. This leads to a two-stage framework suitable to be integrated with existing (or new) AIS and/or SMC samplers. The marginals are adapted in the first stage, while the coupling is adapted in the second stage. We show with experiments the benefits of the proposed methodology, including an application to Bayesian prediction with misspecified models.